US009143406B2

(12) **United States Patent**
Waggener et al.

(10) **Patent No.: US 9,143,406 B2**
(45) **Date of Patent: Sep. 22, 2015**

(54) **APPARATUS, METHOD AND COMPUTER-READABLE STORAGE MEDIUM FOR CALCULATING THROUGHPUT REQUIREMENTS OF A NETWORK**

(75) Inventors: **Michelle A. Waggener**, Colorado Springs, CO (US); **Michael U. Bencheck**, Richardson, TX (US); **Scott Ray Kotrla**, Wylie, TX (US); **Matthew William Turlington**, Richardson, TX (US); **Christopher Nicholas DelRegno**, Rowlett, TX (US); **Richard C. Schell**, Allen, TX (US)

(73) Assignee: **VERIZON PATENT AND LICENSING INC.**, Basking Ridge, NJ (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 758 days.

(21) Appl. No.: **11/873,877**

(22) Filed: **Oct. 17, 2007**

(65) **Prior Publication Data**

US 2009/0103448 A1 Apr. 23, 2009

(51) **Int. Cl.**

H04L 12/24	(2006.01)
H04L 12/54	(2013.01)
H04L 12/801	(2013.01)
H04L 12/927	(2013.01)
H04L 12/28	(2006.01)

(52) **U.S. Cl.**

CPC ***H04L 41/0896*** (2013.01); ***H04L 12/5695*** (2013.01); ***H04L 41/145*** (2013.01); ***H04L 47/10*** (2013.01); ***H04L 47/808*** (2013.01); *H04L 12/2856* (2013.01); *H04L 41/5003* (2013.01); *H04L 41/5009* (2013.01); *H04L 41/5067* (2013.01)

(58) **Field of Classification Search**

USPC 370/230, 235, 252, 352, 395.21; 725/111, 125

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,985,437	B1 *	1/2006	Vogel	370/230
7,127,272	B1 *	10/2006	Almaleh et al.	455/561
7,391,728	B2 *	6/2008	Natarajan et al.	370/235
8,432,849	B2 *	4/2013	Krishnasawamy et al.	370/317
2004/0184535	A1 *	9/2004	Wang	375/240.03
2005/0157728	A1 *	7/2005	Kawano et al.	370/395.21
2006/0030323	A1 *	2/2006	Ode et al.	455/436
2006/0056515	A1 *	3/2006	Kato et al.	375/240.16

(Continued)

OTHER PUBLICATIONS

Indra Widjaja, Alberto Leon-Garcia, "Communication Networks Fundamental Concepts and Key Architectures", 2004, McGraw Hill, ISBN 978-0072463521, pp. 373-375.*

Primary Examiner — Omer S Mian

(57) **ABSTRACT**

An apparatus, method and computer-readable storage medium are provided for calculating an estimated throughput requirement of a network, such as to facilitate selecting the minimum bandwidth, maximum bandwidth or some other appropriate bandwidth for providing network services to a customer accessing that network. In this regard, a method of calculating a throughput requirement of a network includes receiving a bandwidth requirement. The method also includes calculating a data link-layer frame rate requirement based on the bandwidth requirement and a frame size, where the calculated frame rate requirement is limited by a line rate for transmission of data packets across a core network. In addition, the method includes calculating a data link-layer throughput requirement based on the calculated frame rate requirement and the frame size.

12 Claims, 5 Drawing Sheets

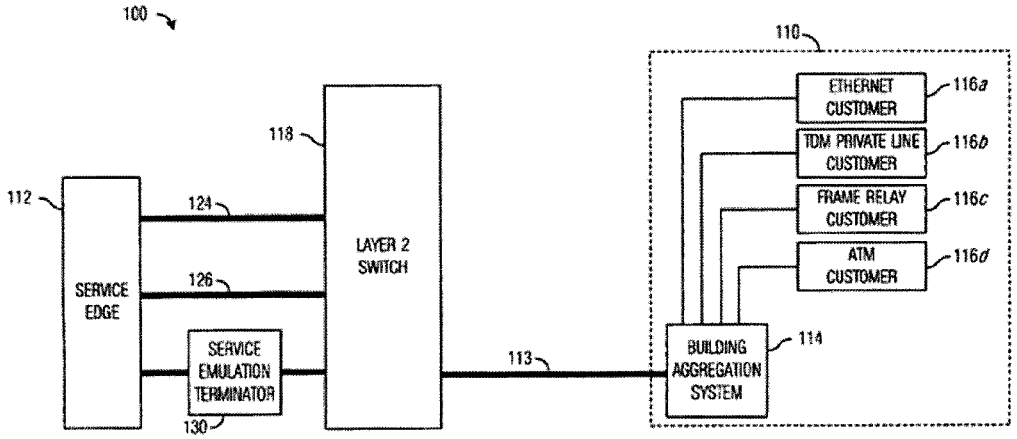

(56) **References Cited**

U.S. PATENT DOCUMENTS

2006/0233100 A1* 10/2006 Luft et al. 370/229
2008/0130790 A1* 6/2008 Forenza et al. 375/299
2008/0209186 A1* 8/2008 Boden et al. 712/225

* cited by examiner

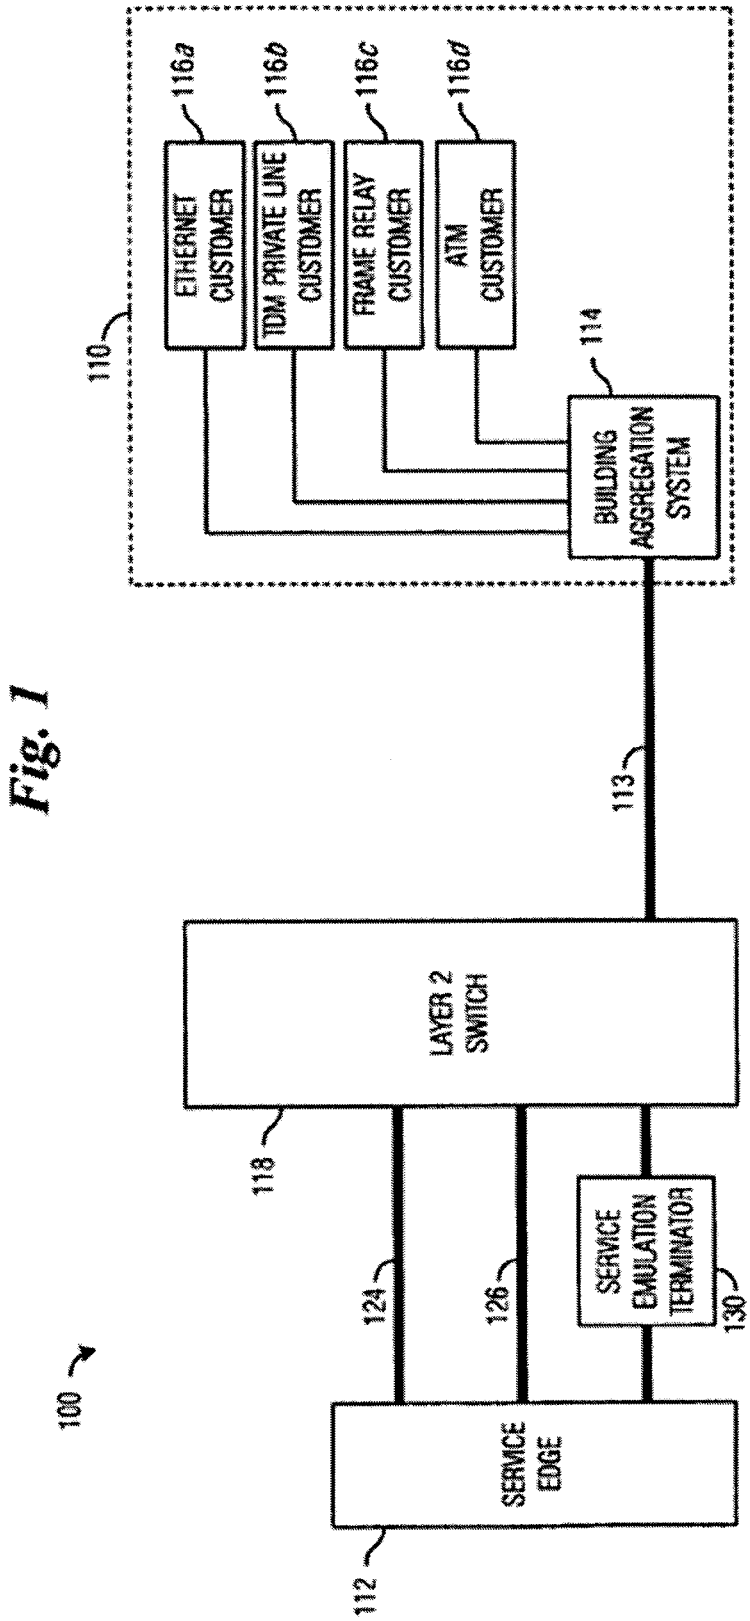
100
110
116a
ETHERNET CUSTOMER
116b
TDM PRIVATE LINE CUSTOMER
116c
FRAME RELAY CUSTOMER
116d
ATM CUSTOMER
114
BUILDING AGGREGATION SYSTEM
113
LAYER 2 SWITCH
118
124
126
SERVICE EMULATION TERMINATOR
130
SERVICE EDGE
112

Fig. 1

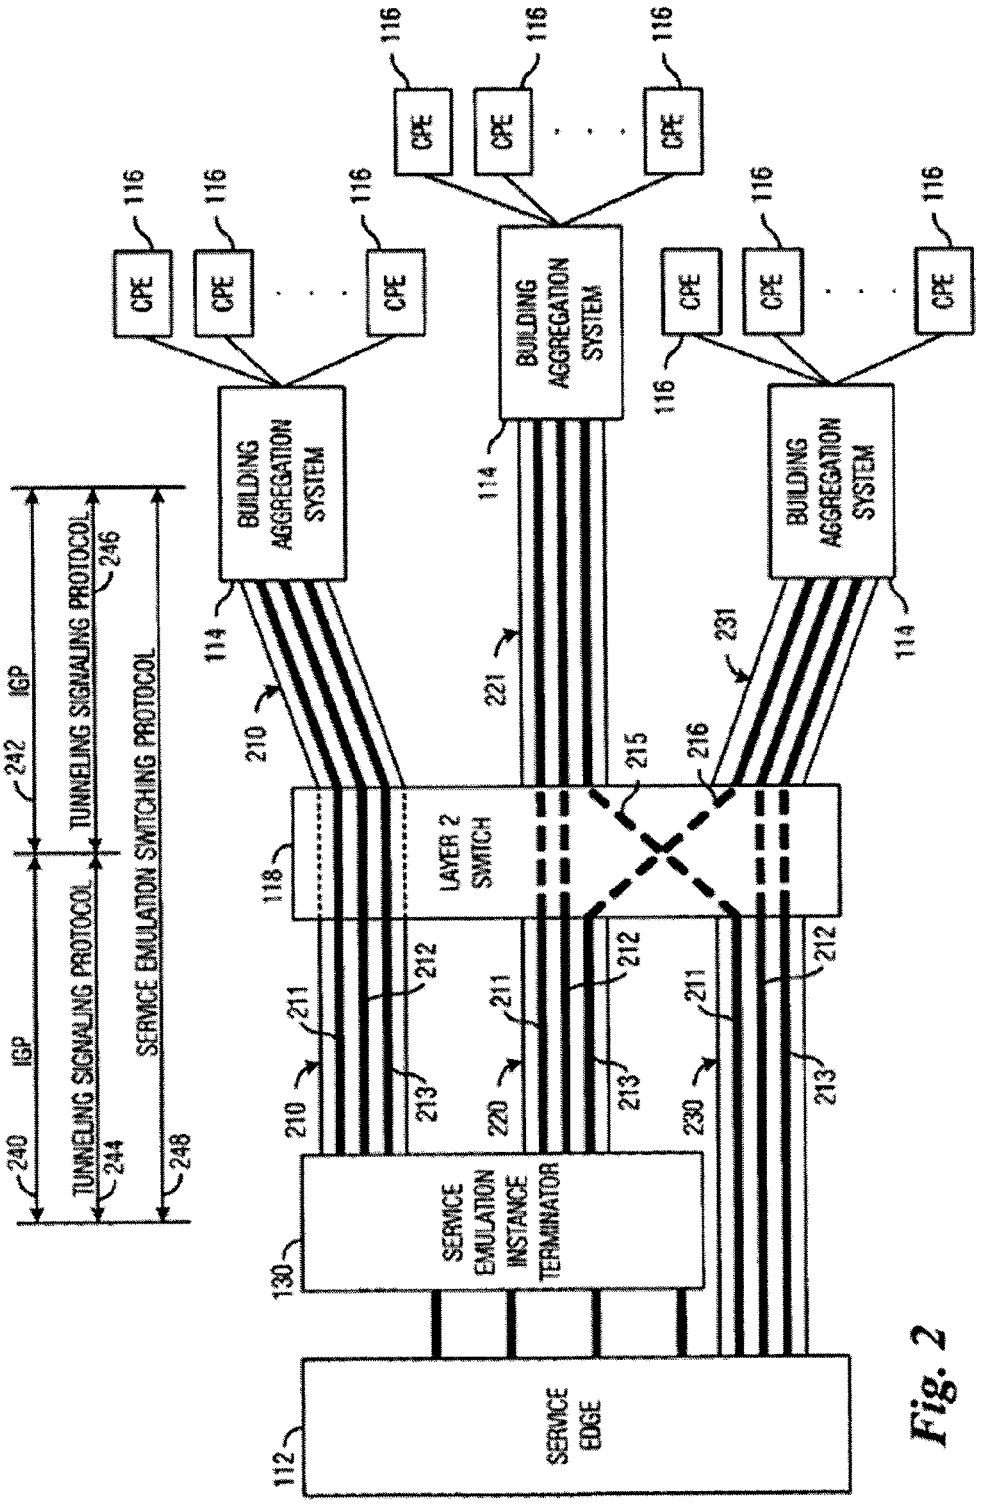
CPE
116
BUILDING AGGREGATION SYSTEM
114
LAYER 2 SWITCH
118
SERVICE EMULATION INSTANCE TERMINATOR
130
SERVICE EDGE
112
210
211
212
213
215
216
220
221
230
231
IGP
TUNNELING SIGNALING PROTOCOL
SERVICE EMULATION SWITCHING PROTOCOL
240
242
244
246
248

Fig. 2

Fig. 3

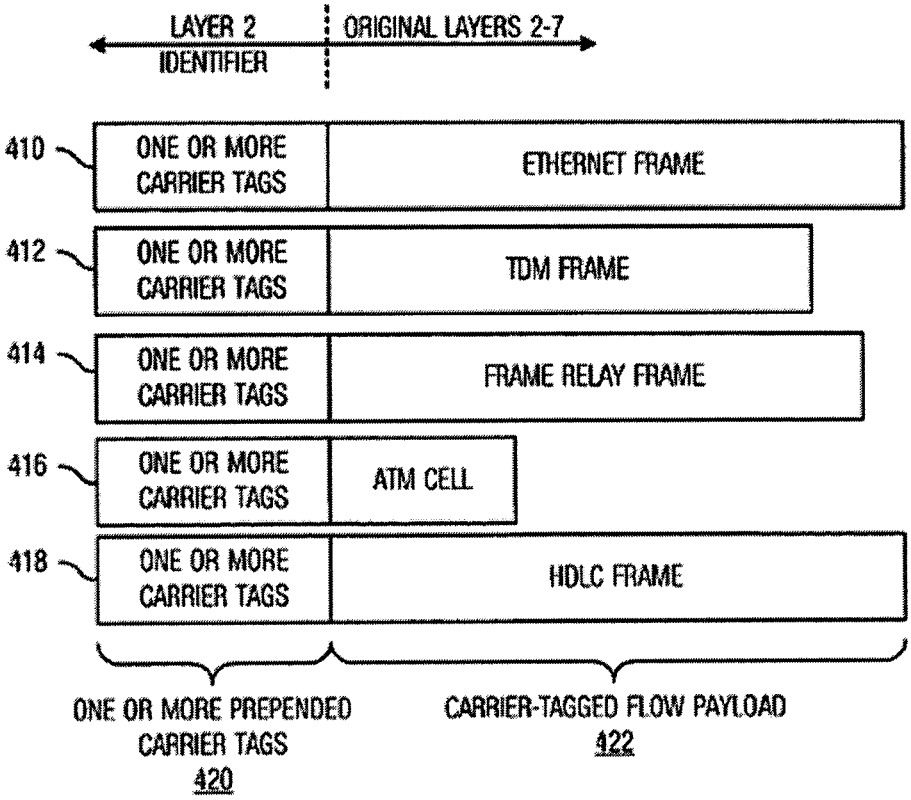
LAYER 2 IDENTIFIER
ORIGINAL LAYERS 2-7
410
ONE OR MORE CARRIER TAGS
ETHERNET FRAME
412
ONE OR MORE CARRIER TAGS
TDM FRAME
414
ONE OR MORE CARRIER TAGS
FRAME RELAY FRAME
416
ONE OR MORE CARRIER TAGS
ATM CELL
418
ONE OR MORE CARRIER TAGS
HDLC FRAME
ONE OR MORE PREPENDED CARRIER TAGS
420
CARRIER-TAGGED FLOW PAYLOAD
422

Fig. 4

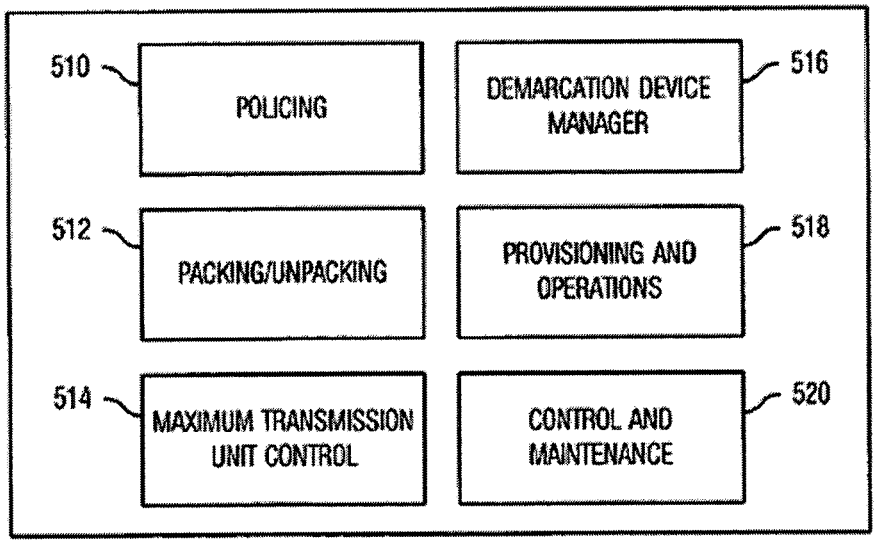
510
POLICING
516
DEMARCATION DEVICE MANAGER
512
PACKING/UNPACKING
518
PROVISIONING AND OPERATIONS
514
MAXIMUM TRANSMISSION UNIT CONTROL
520
CONTROL AND MAINTENANCE

Fig. 5

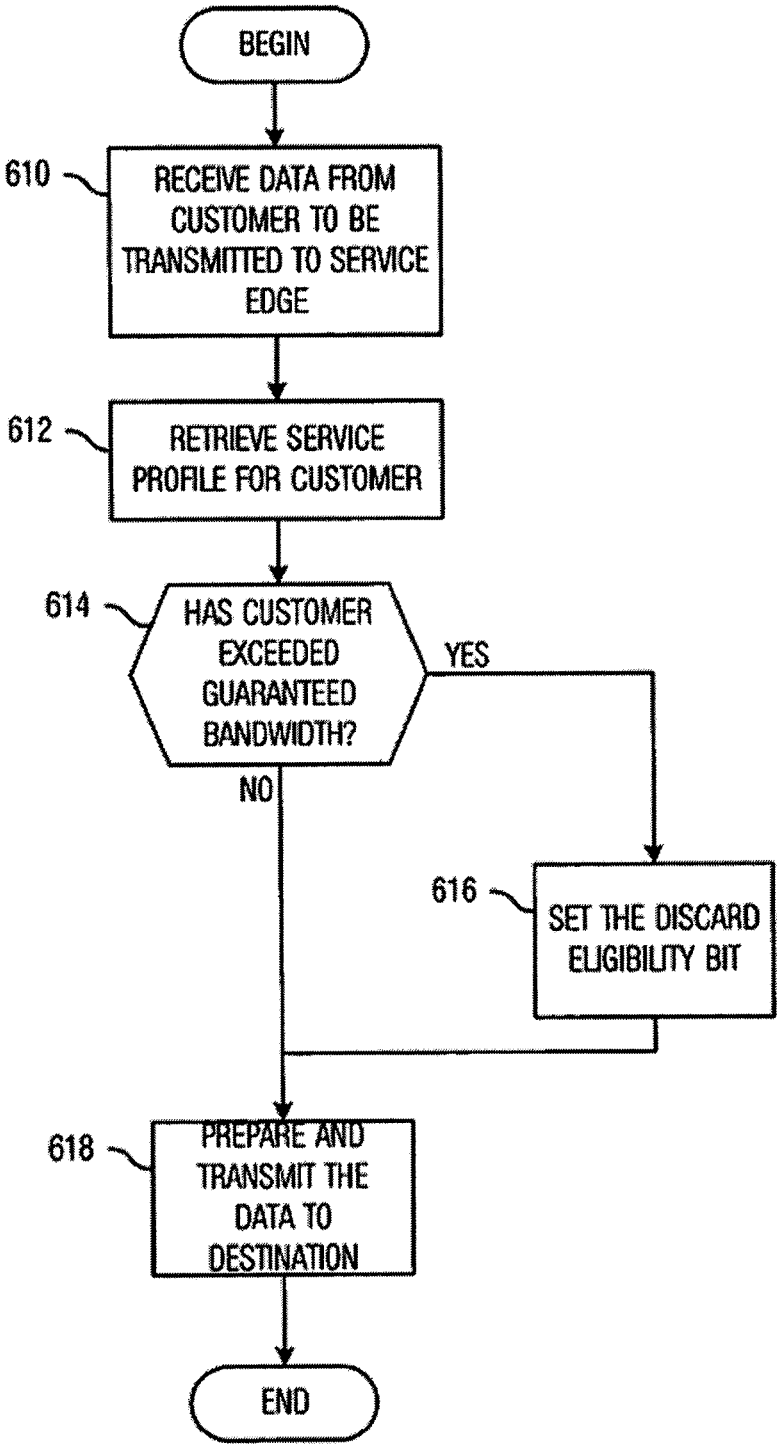
BEGIN
610
RECEIVE DATA FROM CUSTOMER TO BE TRANSMITTED TO SERVICE EDGE
612
RETRIEVE SERVICE PROFILE FOR CUSTOMER
614
HAS CUSTOMER EXCEEDED GUARANTEED BANDWIDTH?
YES
NO
616
SET THE DISCARD ELIGIBILITY BIT
618
PREPARE AND TRANSMIT THE DATA TO DESTINATION
END

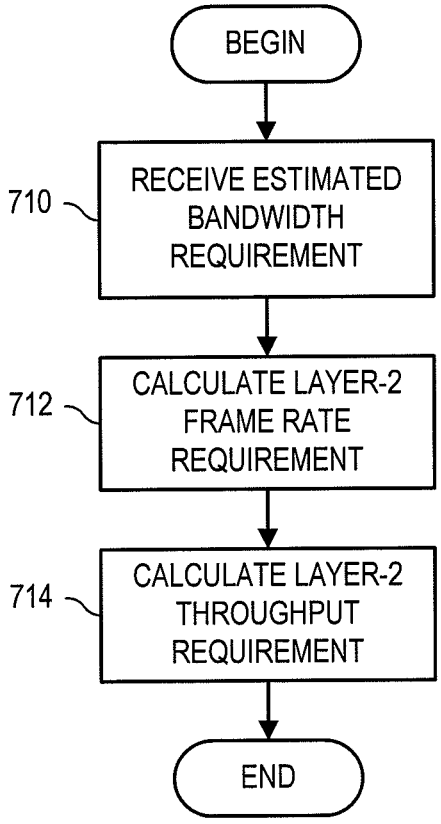
BEGIN
710
RECEIVE ESTIMATED BANDWIDTH REQUIREMENT
712
CALCULATE LAYER-2 FRAME RATE REQUIREMENT
714
CALCULATE LAYER-2 THROUGHPUT REQUIREMENT
END

Fig. 6

APPARATUS, METHOD AND COMPUTER-READABLE STORAGE MEDIUM FOR CALCULATING THROUGHPUT REQUIREMENTS OF A NETWORK

FIELD

An apparatus and method are provided for generally calculating a requirement of a network, and for more particularly calculating a throughput requirement of a network to thereby to facilitate setting up that network for a customer accessing one or more services via that network.

BACKGROUND

A commercial telecommunications network operated by a service provider supports voice and data communications between customer locations and includes an access network and a core network. Generally, customer devices communicatively couple to the access network, which in turn connects to the core network. The access network includes what many people refer to as "the last mile," that is, the connectivity from a customer location, such as an office building, to a point where a service provider has significant facilities, such as a metro hub or a "service edge" at the periphery of the core network. In contrast to the access network, the core network usually provides transport of large aggregate flows over long distances and handles the selective routing of each customer's voice and data traffic to other locations served by the network. The access network generally comprises a series of switches, aggregators, multiplexers, demultiplexers, routers, hubs, and the like which collectively serve to provide connectivity between customers' equipment and the core network.

Customer sites in the vicinity of a service provider's edge, or an intermediate hub that provides connection to the service edge, must be connected to the service edge via some form of access circuit. Traditionally, it has been more practical for a core network service provider to operate a few strategically placed facilities to serve a large number of customers in a metropolitan area rather than to extend the service provider's core network to every physical location where customers may reside. Providing access services between a customer's location and a metro hub or a service edge may involve installing electrical or optical cables between the service provider and the customer site. In some cases, the service provider installs and owns this access link connected directly to the customer location. More often, however, the existing facilities of a local telephone carrier are leased to provide this connectivity. The well-established local telephone facilities provide at least twisted-pair subscriber loop connectivity to virtually every potential customer location in a metropolitan area. In the case of larger business locations and multi-tenant commercial sites, local telephone facilities typically comprise a large quantity of telephone wires or broadband access to the sites.

The services required by customers, residential or business, vary greatly in the type of access services, bandwidth, quality of service (QoS), type of legacy equipment, and the like. Types of services typically include frame relay services, asynchronous transfer mode (ATM) services, broadband services, point-to-point private line services, voice services, and the like. Accordingly, the access service provider must be prepared to provide many types of services to customers of various sizes and needs.

Furthermore, the access service provider must be capable of meeting the customers' current and future needs in terms of bandwidth, throughput, QoS, and the like. For example, a given customer may start with relatively small bandwidth needs yet expand to needing high-bandwidth connections such as a SONET OC-3 or OC-12 connection. Additionally, customers may require ATM services and frame relay services to support legacy systems while implementing newer applications and communications that are based on Ethernet.

Typically, the access service provider provisions and handles each of these services separately and does not attempt to aggregate the traffic from various types of services. Treating the traffic separately in this manner, however, can be expensive. Typically, the access network provides transport, aggregation, grooming, and switching for each of these types of services independently, which in turn requires the access service provider to provision each of these services separately. Generally, each type of service utilizes different interface and framing standards, and in particular, each type of service typically utilizes a different set of protocols. As a result, current access network elements must be equipped to interface with and operate upon flows for each type of service the elements are expected to handle. Each network element in an access network presently must deal with the particular format, addressing and protocol aspects of each type of access communication service it supports. This makes for costly and complex network elements and interferes with having flexibility to accommodate rapid shifts in resources allocated to different flows or different service types and to accommodate adoption of new service types.

Although techniques have been developed for providing services to customers in a cost effective manner that allows the service provider to provide various types of services to one or more customers, it is generally desirable to improve upon existing techniques.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1** is a network diagram of an access network embodying features of exemplary embodiments;

FIG. **2** is logical view of the use of service emulation instances in the access network in accordance with one exemplary embodiment;

FIG. **3** illustrates examples of data messages or frames that may be transmitted into the access network, or received from the access network, in accordance with an exemplary embodiment;

FIG. **4** is a functional block diagram of a building aggregation system in accordance with one exemplary embodiment;

FIG. **5** is a flowchart including various steps in a method of marking or dropping of customer packets that may be performed by a policing function of the building aggregation system, in accordance with an exemplary embodiment; and

FIG. **6** is a flowchart including various steps in a method of calculating a throughput requirement of a network, in accordance with an exemplary embodiment.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

Exemplary embodiments are described hereinafter with reference to the accompanying drawings, in which exemplary embodiments and examples are shown. Like numbers refer to like elements throughout.

Exemplary embodiments will be described in a specific context, namely, providing aggregation services to one or more customers in an office building communicatively coupled to a telecommunications network via a TDM (time division multiplexing) communications link, for example a

DS3. The exemplary embodiments may also be applied, however, to other types of devices, networks, communications links, and the like. Furthermore, while specific network configurations are illustrated and discussed herein, it is noted that network configurations may vary to include fewer or additional elements, such as routers, gateways, bridges, ATM (asynchronous transfer mode) switches, frame relay switches, firewalls, and the like. The illustrated embodiments are provided for illustrative purposes only and are provided only to aid in the explanation and understanding of the concepts. Aspects are equally applicable to many types and configurations of networks and communications protocols.

It is further noted that, unless indicated otherwise, functions described herein may be performed in hardware, software, firmware or some combination thereof. In one exemplary embodiment, however, the functions may be performed by a processor such as a computer or an electronic data processor in accordance with code such as computer program code, software, and/or integrated circuits that are coded to perform such functions, unless otherwise indicated. It should be understood, however, that some of the functions described herein may not be performed in hardware, software, firmware or some combination thereof. Instead, these functions may be performed by one or more persons or users who may be users of such hardware, software, firmware or combination. Thus, an apparatus, may be shown, described and claimed as including a processor configured perform and/or facilitate performance of a method. In such instances, the processor may perform one or more steps of that method in concert with one or more persons acting to perform other steps of that method.

FIG. **1** is a network diagram in accordance with an exemplary embodiment. FIG. **1** illustrates an access network **100** to which a customer site **110** may be coupled, such as to access the services of a service edge **112**. The customer site include, for example, an office building housing a large number of users, such as employees of one or more business enterprises. The users or businesses may subscribe to the services of a service provider whose point of presence is the service edge.

Generally, the service edge **112** may represent one or more access points to the service provider's network, which may comprise one or more core networks (not shown). A core network may include, for example, a system of TDM switches, such as a network of Class 3 telephone switches. A core network may also include an ATM and/or a frame relay network covering much the same geographical territory as the TDM network. Moreover, a network of IP (Internet Protocol) routers may also be supported in a core network. While each of one or more of these networks may overlap or cover much the same geographical territory, the respective networks may be designed to efficiently carry particular types of traffic or exhibit particular properties that are amenable to certain types of traffic. Although this "multi-planar" network situation may frequently be encountered, it should be understood that exemplary embodiments may be equally applicable to a converged core network where native layer-2 handoff at the service edge may be desirable. The service edge is illustrated as a single network element for illustrative purposes only, but may actually include multiple network elements or multiple access interfaces having different capabilities.

By way of example, sources of, or customer premises equipment (CPE) for, different types of communications are depicted within a customer site **110**. These sources may include an Ethernet customer **116***a* coupled to a building aggregation system (BAS) **114** over any form of connectivity amenable to Ethernet traffic, such as a 100BaseT, Gigabit Ethernet (GbE) or DSL (digital subscriber line) connection. Another source of traffic may be private line customer **116***b*, which may be coupled to the building aggregation system via DS1 line. Source **116***c* represents frame relay customers having their frame relay traffic carried over TDM facilities such as DS1 lines to the building aggregation system. ATM customer **116***d* represents ATM customers having their ATM cell traffic carried over TDM facilities such as DS1 lines to the building aggregation system. Other types of connections may be used as required to support specific customers' needs. Each of one or more CPEs **116***a*-**116***d* may comprise one or more devices. For example, the Ethernet customer **116***a* may include a router communicatively coupled to other routers, hubs, user workstations, servers, or the like. The CPEs **116***a*-**116***d* may be collectively referred to as CPE **116**.

As noted above, customers within a building may require different types of access, or a single customer may require different types of access for different services. As will be discussed in greater detail below, the building aggregation system **114** may provide an interface to one or more CPEs, which may be operating in accordance with one or more communications protocols; and may aggregate the traffic for transmission to the service edge **112**. In this manner, economies of scale may be realized by combining traffic from a plurality of less-than-fully utilized communications channels. A building aggregation system may serve multiple buildings within a reasonable proximity, such as a corporate, institutional campus or any other collection of sites where it is feasible.

Other components, such as demarcation devices, repeaters, amplifiers, and the like, may be communicatively coupled between the building aggregation system **114** and each of one or more CPEs **116***a*-**116***d*. A demarcation device, commonly used when providing Ethernet services (such as that provided to an Ethernet customer, represented by CPE **116***a*), is the point at which the customer connects to the resources of the access network **100**. Further, additional repeaters (not shown) or amplifiers (not shown) may be required based upon, for example, the length of the wire runs.

On the network side, the building aggregation system **114** may be communicatively coupled to one or more hubs or switches, represented by switch **118**, to provide connectivity between the customer location **110** and the service edge **112**. The communications link **113** between the building aggregation system and the switch may include, for example, a TDM communications link, such as a DS3 or SONET OC-n. In accordance with one exemplary embodiment, these TDM links utilize a protocol such as X.86, GFP (Generic Framing Protocol), PPP (Point-to-Point Protocol), or the like for accomplishing packet data transport over TDM. The communications path between the customer location and the service edge is illustrated as a simple two-hop connection for illustrative purposes only. It should be understood, however, that the communications path between the customer location and the service edge may contain additional or fewer hops, and may include different paths for traffic in either direction between a service edge and a customer site. The customer location, through the building aggregation system, may be coupled to service edge through a network of switches and other equipment and facilities.

Additional network elements may be positioned between the building aggregation system **114** and the switch **118**. For example, the building aggregation system may be configured to accept a DS3 communications link as discussed above, but the communications link from the access network to the building **110** may comprise a larger communications link, such as an OC12 or OC48 optical link. In such instances, which are common in an "on-network" environment wherein the access network is owned by the service provider, an add/

drop multiplexer (ADM) may be utilized to separate the DS3 traffic from and interject the DS3 traffic onto the larger OC12 or OC48 link. In an "off-network" environment, on the other hand, a smaller DS3 link may be leased directly from another party, such as a local telephone company or other service provider.

In accordance with one exemplary embodiment, the layer-2 switch **118** may provide switching and routing of traffic based upon information applied to the traffic and without having to examine the content of the customer data in the traffic. The information applied to the traffic may correspond roughly to Layer 2 or the "data link layer" of the OSI Reference Model. The layer-2 switch may be coupled to a large number of customer sites **110** and the building aggregation systems **114** to perform an intermediate aggregation and distribution function within the access network **100**. In some instances, the layer-2 switch may also be coupled directly to some or all of the CPE **116**. An example of a layer-2 switch suitable for use with exemplary embodiments is disclosed in U.S. patent application Ser. No. 10/858,517, entitled: System and Method for Providing a Multiple-Protocol Crossconnect, the content of which is hereby incorporated by reference in its entirety.

The building aggregation system **114** may apply information to the traffic that has significance for affecting the handling of carrier-tagged communications within the access network and may be interpreted and acted upon by layer-2 switch **118** or other elements. The building aggregation system may also receive communications bearing this information and may route the communications to specific customers in response to the information.

In accordance with exemplary embodiments, the building aggregation system **114** may be equipped or otherwise configured to serve as one end of a plurality of carrier-tagged flows. In this regard, a carrier-tagged flow represents a logical communications channel or flow established to carry carrier-tagged communications between two or more parties, or two or more points served by a communications system. The carrier-tagged communications may include, for example, voice, data, audio, video or any other type of communications.

A carrier-tagged flow may be implemented using a service emulation instance, such as a pseudowire as described IETF (Internet Engineering Task Force) Request for Comments document RFC 3985, entitled: Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture (March 2005), the content of which is hereby incorporated by reference in its entirety. This technology allows a packet-switched network to emulate other types of packet or TDM transport services. For example, a pseudowire may be implemented in an Ethernet network, yet may provide transport of communications that mimics the attributes and performance of common data link protocols, such as ATM, frame relay, as well as SONET/SDH or DSn signals. An Ethernet-based pseudowire may employ variable length packets even when carrying fixed-length cells or frames, such as 53-byte ATM cells.

A pseudowire may be implemented along a tunnel implemented in a packet-switched network. Some types of tunnels that may be suitable for carrying pseudowires, or other types of communications that may be employed in conjunction with exemplary embodiments, include Label Switched Paths (LSPs) according to the Multi-Protocol Label Switching (MPLS) protocol, Layer-2 Tunneling Protocol (L2TP) tunnels, IPsec tunnels or the like.

Another example of a technique suitable for implementing a carrier-tagged flow is a logical networking tagged flow, such as virtual local-area network (VLAN) communications or the like. One exemplary technique for achieving VLAN logical sub-networking is described in IEEE Standard 802.1Q. Briefly, a VLAN may enable designating and acting upon data packets in a manner that makes multiple LAN communication flows carried over a commonly shared communication path appear to be partitioned from one another as if traveling over separate, dedicated LAN connections. In accordance with an exemplary embodiment, a VLAN tagging approach may also be used for carrier-tagging of flows.

In accordance with exemplary embodiments, carrier VLAN tags having significance for routing and processing in the access network may be used to encapsulate and tag customer flows. As customer flows are encapsulated and/or tagged, they may or may not already include additional, imbedded VLAN tags having significance within the customer's virtual network in accordance with conventional IEEE 802.1Q usage. In accordance with exemplary embodiments, the VLAN tagging approach may be reused for carrier-tagging purposes and may be locally significant on any port, with tag values possibly being replaced on a hop-by-hop basis.

In accordance with exemplary embodiments, carrier tags applied to traffic to support handling of flows through an access network, whether in the form of tunnel labels, pseudowire labels, VLAN tags or the like, may be 'stackable' to any depth to support efficient flow management in the context of hierarchical aggregation and distribution between service edge(s) and customer locations.

In addition to supporting access communications that enable carrier tag stacking, a building aggregation system **114** in accordance with exemplary embodiments may serve several roles in the context of a service-agnostic access network. The building aggregation system may adapt a wide variety of customer traffic to be carried in the form of carrier-tagged flows. Where a packet switched network is used as an underlying access transport, the use of carrier-tag switching may enable the efficiencies and flexibility of packetization to be realized in an access network without burdening the access network elements with the specific protocols or addressing used in the carried traffic.

To coordinate the assigning of carrier tags to traffic flows, the building aggregation system **114** may participate in label resolution protocols with other elements, and may receive provisioning instructions from remote management systems. To support the deployment of service emulation, such as pseudowire technology, in an access network to achieve service flexibility, the building aggregation system may serve as a terminal end for a large number of service emulation instances of different emulated types and may provide mapping and forwarding between customer flows and access network paths as identified by tunnel labels or service emulation instance mapping identifiers. The building aggregation system may also implement QoS (quality of service) functions for all types of flows, augmenting similar measures that may be performed at a service edge. Other possible roles and functions of the building aggregation system are possible.

The building aggregation system **114** may be used in a carrier-tag oriented access network wherein each of one or more carrier-tagged flows is identified by a carrier tag having a particular tag value. For example, a carrier-tagged flow implemented as a service emulation instance is identified by a service emulation instance mapping identifier. In the case in which a pseudowire is used as a service emulation instance, the service emulation instance mapping identifier may take the form of a pseudowire label. Carrier tags may be locally significant on any port and the tags can be swapped on a hop-by-hop basis as needed to provide a large number of flows using the finite number of identifier values that are

available (approximately one million in the case of pseudowire labels). In this manner, switching in the access network can be simplified by encapsulating traffic in carrier-tagged flows and by interpreting and manipulating the corresponding carrier tags.

An access network **100** in accordance with exemplary embodiments may transparently support a mixture of flow types and customer content, including any customer-specific addressing or virtual networking fields imbedded in the customer content. The pseudowire architecture, as described in documents promulgated by the IETF, may provide one example of an approach involving encapsulation and labeling of traffic that may be adapted for use as a carrier-tagged flow. It should be noted, however, that other protocols may be used, and exemplary embodiments may be implemented with other types of protocols and physical connections.

The building aggregation system **114** may couple traffic of various types, such as traffic from the CPE **116***a*-**116***d*, onto the appropriate corresponding carrier-tagged flows established for reaching the service edge. Where service emulation instances are used as carrier-tagged flows, a service emulation instance terminator **130** may serve as the other end of a number of service emulation instances which have originated at one or more building aggregation systems and passed through layer-2 switches **118**. The service emulation instance terminator may switch or route traffic from service emulation instances to a corresponding port and/or flow communicatively coupled to the service edge **112**. The building aggregation system, layer-2 switch, service emulation instance terminator and the communications links therebetween may coordinate to simultaneously function as any of the various data-link layer transport types that may be required by customers, including ATM, frame relay, TDM, Ethernet/IP, and the like.

Alternatively, a service edge **112** may incorporate the functions of a service emulation instance terminator **130** or may otherwise be capable of directly accepting and processing carrier-tagged flows. In this case, a service edge may be coupled more or less directly to the layer-2 switch **118** and the communications to and from the service edge may bear flow-identifying carrier tags in the form of pseudowire labels, tunnel labels, VLAN tags or the like. The service emulation instance terminator may nevertheless be useful in situations where an existing or legacy service edge facility lacks the ability to handle carrier-tagged access communications, service emulation instances or, more specifically, pseudowires. As mentioned earlier, the service edge may actually represent several separate access points, perhaps to different types of core networks. Some access points within the service edge may be amenable to carrier-tagged flows whereas others may not be. Links **124** and **126** may represent links to TDM-capable ports on the service edge from TDM ports on the layer-2 switch. It is also possible that, for example, one or both of these links may represent packetized data links or may represent a service edge that is able to accept carrier-tagged flows, such as carrier-tagged pseudowires, directly without requiring the service emulation instance terminator. A service emulation instance terminator suitable for use with exemplary embodiments is disclosed in U.S. patent application Ser. No. 10/858,491, entitled: Apparatus and Method for Terminating Service Emulation Instances, the content of which is hereby incorporated by reference in its entirety.

In an exemplary embodiment, Ethernet may be utilized as the layer-2 protocol over which carrier-tagged communications are transmitted. The application of Ethernet in the access network can be based on TDM encapsulation using X.86 or GFP, such as in the case of Ethernet over SONET (EoS). While Ethernet is desirable because it supports variable length packets or frames, other protocols or frame formats may be used for the transport and processing of access communications.

In an implementation using service emulation instances, the building aggregation system **114** may apply a unique service emulation instance mapping identifier to each of one or more flows from the CPE **116***a*-**116***d*, and transmits the frames or packets bearing the traffic and service emulation instance mapping identifiers to the layer-2 switch **118**. Similarly, the building aggregation system may receive data associated with a service emulation instance identifier from the layer-2 switch and convert the data to a format compatible with the corresponding CPE **116**.

As another example, the building aggregation system **114** may receive Ethernet traffic from Ethernet customer **116***a* via the building "riser." In such instances, the building aggregation system may receive this traffic along a port that is known to correspond to Ethernet customer **116***a* and maintain an association between the customer's port and Ethernet traffic stream and a corresponding carrier-tagged flow. Likewise, at the other end of the carrier-tagged flow, the service emulation instance terminator **130**, the layer-2 switch **118**, or some other network element may deliver the customer's traffic to the service edge **112** and may coordinate with the service edge, such as by mapping of port numbers or directing of flows, to at least partially ensure that the network identifies the customer's traffic as such and appropriately handles the traffic.

In accordance with one exemplary embodiment, when a device on a customer LAN addresses Ethernet frames to the service edge **112**, the destination MAC address field in the Ethernet frames may correspond to that of the service edge itself or some other remote device port rather than being associated with any of the access network elements. CPE **116** may be provisioned to resolve the address of the service edge such that communications intended to be sent to the service edge will arrive at a port of the building aggregation system. The building aggregation system **114** does not need to act as a layer-2 termination point as it appears to the CPE, so no MAC learning behavior or other typical LAN bridge behavior is required of the building aggregation system.

To establish or modify a customer's carrier-tagged flow between the CPE **116** and the service edge **112**, the customer may indicate to the network service provider the desire to establish communications in a particular manner. This request may be submitted either manually or automatically through a user network interface (UNI). The establishment of communications through the access network **100** shown may originate in a variety of ways. To coordinate fulfillment of an access communications request, a network management system, provisioning function or the like, may dispatch provisioning and configuration instructions to the building aggregation system **114**, layer-2 switch **118**, service emulation instance terminator **130** or other network elements. To varying degrees, these elements may perform some functions autonomously or may coordinate with one another to fulfill service requests.

For convenience, the operation of exemplary embodiments discussed herein maybe described in terms of traffic flowing from the CPE **116** to the service edge **112**. It should be noted, however, that the same techniques discussed herein also apply to traffic leaving the service edge and being distributed to an appropriate customer endpoint. Every element may serve a complementary role related to the direction of flow. It is also worth noting that some traffic through an access network **100** may be from one customer location to another in a given vicinity, and may not necessarily be destined for the

service edge. Many of the techniques described for traffic between a customer location and a service edge would be applicable to this situation as well.

Reference is now made to FIG. **2**, which illustrates a logical view of example carrier-tagged flows through an access network in accordance with an exemplary embodiment. Customer sites (represented by CPE **116** coupled through building aggregation systems **114**) are shown to be coupled, through various access network resources, to the edge of a service provider's network represented by the service edge **112**, and optionally involving service emulation instance terminator **130**. The layer-2 switch **118** is shown as an intermediary and may participate in grooming, aggregating and directing communications traffic in the access network, as well as performing crossover switching between TDM ports and packet-oriented ports. For clarity, FIG. **2** illustrates two-hop paths, although it is possible that there are some intervening transmission elements or other layer-2 switches along the access coupling.

In FIG. **2**, tunnels in the form of label switched paths (LSPs) **210**, **220**, **221**, **230** and **231** are shown to have been established between various building aggregation systems **114** and the service edge **112** (optionally through terminator **130**). Each of one or more LSPs corresponds to a pathway or a tunnel for carrying traffic from the building aggregation system **114** to the service edge **112**, or vice versa. It should be noted that FIG. **2** is provided as a logical view and that various physical implementations may be used. For example, each of one or more label switched paths may be transported, routed or switched over different physical links or the same physical link.

Each LSP is shown to be carrying a respective set of carrier-tagged flows **211**, **212** and **213**. In some implementations, such as when service emulation instances are used, each of one or more carrier-tagged flows may emulate a specific type of service, such as an ATM, frame relay, Ethernet or TDM/SONET service. Regardless of emulated service type, the traffic belonging to these carrier-tagged flows may be made to travel along an LSP or tunnel based upon a carrier tag applied to the traffic and a mutual understanding among the network elements as to how to handle traffic having a specific tag value.

Where service emulation instances are used to implement carrier-tagged paths, the traffic may be identified and routed through the access network **100** based on a service emulation instance mapping identifier that serves as the carrier tag (or as one part of a composite carrier tag). More specifically, where pseudowires are used as service emulation instances, a pseudowire label in each of one or more traffic packets or frames may serve as a carrier tag component.

In an implementation where logical networking tagged flows are used, an outer VLAN tag may serve as the carrier tag, or a portion thereof. A tunnel label, such as an LSP label, may also be applied to the traffic and may be used, solely or in conjunction with other fields, as the label upon which the traffic is routed through access network elements. In accordance with one exemplary embodiment, any of these possible carrier tags may be present multiple times or may be stacked in any combination with one another within the traffic packets or frames. This may be done to support various flow management arrangements that involve nesting of tunnels or flows or otherwise implementing stages of aggregation in the access network. Examples of how nested tunnels or paths may be accomplished and managed are disclosed in U.S. patent application Ser. No. 10/858,525, entitled: System and Method for Managing Communications in an Access Network, the content of which is hereby incorporated by reference in its entirety.

Each of one or more carrier-tagged flows **211**, **212**, **213** may carry multiple customer-specified "subflows" in a manner that is transparent to the access network. That is, the customer traffic may optionally contain additional imbedded VLAN tags having significance within the customer's virtual private network in accordance with conventional 802.1Q usage. Any VLAN tags or layer-2 VPN addressing fields that are present within the customer traffic may be encapsulated using carrier VLAN tags having significance for routing and processing in the access network. Consequently, a traffic frame or packet in accordance with exemplary embodiments may comprise some VLAN tag fields that are controlled by a customer, or have significance within the customer's private network, in addition to VLAN tags that serve as carrier tags having significance to the access network. Compared to any customer-imposed VLAN tags appearing in the traffic, the carrier VLAN tags may be derived from the operation of entirely different protocols among different elements than the customer tags. Likewise, where other forms of carrier tags are employed, such as pseudowire labels, the carrier tags may resemble customer labels or tags appearing in customer traffic, but may have different significance and may be derived from different protocols than the customer information. In accordance with exemplary embodiments, the access network **100** may be primarily concerned with the outermost labels or carrier tags that have been applied to the traffic for access network purposes, such as tunnel labels or service emulation instance mapping identifiers applied to the traffic. In accordance with exemplary embodiments, a VLAN tag format may be applied for carrier-tagging purposes and may be locally significant on any port, with tag values possibly being replaced on a hop-by-hop basis.

Label switched path **210** represents an exemplary embodiment in which the carrier-tagged flows are routed through the layer-2 switch **118** on the basis of a tunnel label. In other words, each of one or more units of traffic may be tagged with a tunnel label for use by access network elements to determine how to process, and where to send, the traffic. In such instances, a tunnel may be established from the building aggregation system **114** to the service emulation instance terminator **130**. Each of one or more flows within the tunnel identified by the tunnel label, e.g., label switched path **210**, may be routed or switched in the same manner, as illustrated by the dotted label switched path line and the solid-lined flows passing through the layer-2 switch. With this approach, the layer-2 switch may efficiently switch traffic among its ports by observing and acting solely or primarily upon this tunnel label present in the traffic. The layer-2 switch may not have to read and act upon a service emulation identifier, such as a pseudowire label, in order to properly route traffic of this nature.

In an alternative embodiment, the carrier-tagged flows **211**, **212**, **213** may be service emulation instances and each service emulation instance may be routed or switched based upon a service emulation instance mapping identifier. As shown, for example, the label switched paths **221** and **231** may be established between the various building aggregation systems **114** and the layer-2 switch **118**. The other LSPs **220** and **230**, then, may be separately established between the layer-2 switch and the service edge **112**. This provides the option of switching individual flows within the layer-2 switch. Switching within the layer-2 switch may be based upon a service emulation instance mapping identifier present in the traffic.

As in the above example where tunnels directly connect network elements without passing through other switching or terminating elements, tunnel labels may be optional due to the point-to-point nature of the communication. Nevertheless, a tunnel label field may be included to convey useful information, possibly aside from tunnel identification.

To illustrate flow switching within the layer-2 switch **118**, certain of the flows **211**, **212**, and **213** are illustrated as remaining together within each of the labeled switched paths. Upon reaching a switching point, such as the layer-2 switch, at the terminus of a tunnel, such as LSP **221**, each of one or more flows through the access network **100** may be switched independently based upon, among other things, the type of service being provided, the requested service edge, one or more aspects of the traffic, and the like. As indicated by the dotted lines **215** and **216**, each of one or more of the carrier-tagged flows within label switched paths **220**, **221**, **230** and **231** may be routed or switched independently of each other as they pass through the layer-2 switch, where the identification and switching of each of one or more flows may be based upon a flow identifier such as a pseudowire label. In contrast, the flows within LSP **210** may not be switched at the layer-2 switch because the tunnel extends between the building aggregation system **114** and the service emulation instance terminator **130** and, consequently, the layer-2 switch may only observe and act upon the tunnel label. A tunnel such as LSP **210** may have been established to expressly avoid switching of individual flows through the layer-2 switch. In practice, either approach, i.e., "tunneling to" or "tunneling through" the layer-2 switch, may be desirable under various circumstances.

As depicted by reference numeral **248**, a label selection or service emulation switching protocol, such as the Label Distribution Protocol (LDP), may be exercised among the endpoints of a tunnel, a carrier-tagged path, a service emulation instance, a pseudowire, or the like in order to assure agreement among network elements on how traffic will be identified within the tunnel or path. In many cases, these endpoints may be the building aggregation system **114** and the service edge **112** or service emulation instance terminator **130**.

Reference numerals **240** and **242** represent how routes may be chosen between the building aggregation system **114** and the layer-2 switch **118**, and between the layer-2 switch and the service emulation instance terminator **130** or service edge **112**. Identifying and selecting the appropriate paths through the access network may be accomplished using an interior gateway protocol (IGP) such as the Open Shortest Path First-Traffic Engineered (OSPF-TE) approach as described in IETF RFCs 2328, 2676, et al., the content of all of which is hereby incorporated by reference in its entirety. Other routing protocols are known and may be used.

Reference numerals **244** and **246** may indicate that a tunneling signaling protocol, such as the Resource Reservation Protocol (RSVP), may also be used in conjunction with other techniques during establishment of the label switched paths so that the elements involved along the path commit to allocating a specific quantity of bandwidth and other resources to support the requested flow and its performance requirements. Alternatively, it may be possible to establish static LSPs wherein little or no signaling is required. MPLS is described in documents IETF RFCs 3031, 2702, et al. maintained by the IETF, the content of all of which is hereby incorporated by reference in its entirety. Related to the negotiation of labels that are used in MPLS, the label distribution protocol (LDP) is described in IETF Request for Comments document RFC 3036, entitled: LDP Specification (January 2001), the content of which is hereby incorporated by reference in its entirety. The label distribution protocol is also discussed in IETF Request for Comments document RFC 4447, entitled: Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) (April 2006), the content of which is also hereby incorporated by reference in its entirety. The use of RSVP, MPLS and LDP are shown by way of example only and should not be construed as limiting the ways in which exemplary embodiments may be implemented.

The directionality of the traffic may have implications for the establishment of tunnels, service emulation instances, pseudowires, logical networking tagged paths, etc. For example, where an RSVP/LDP mechanism is used to establish label switched paths, a bi-directional link may require initiating the formation of a tunnel in one direction, originating at the building aggregation system **114**, and forming the corresponding tunnel in the reverse direction by originating an RSVP request from the service emulation instance terminator **130**. These tunnels may be independently formed, may have different QoS requirements, and may take different routes between the building aggregation system **114** and the service emulation instance terminator.

FIG. **3** illustrates examples of data messages or frames that may be transmitted into the access network **100**, or received from the access network, by the building aggregation system **114** in accordance with an exemplary embodiment. Each of one or more messages **410-418** may have two portions: a carrier-tagged flow payload **422** and one or more prepended carrier tags **420**. The carrier-tagged flow payload **422** may represent, for example, the information as it is received from customer premise equipment at the customer site. The different types of messages shown correspond to various formats associated with a particular CPE interface, such as, for example, an Ethernet frame message **410** that may be associated with the Ethernet interface module **310**, a TDM frame message **412** that may be associated with the private line interface module **312** or other TDM CPE, a frame relay frame message **414** that may be associated with the frame relay interface module **314**, an ATM cell message **416** that may be associated with the ATM interface module **316**, or the like. Other messages, such as a high-level data link control (HDLC) frame **418**, an ATM application adaptation layer-5 (AAL5) protocol data unit (PDU) or the like, may also be used. In general, the messages may carry various types of customer data corresponding to layers 2-7 of the OSI Reference Model.

As FIG. **3** shows, each of one or more message types may be tagged and processed in a uniform manner by the addition of one or more carrier tags. FIG. **3** reflects the format of composite messages that may be sent between a building aggregation system **114**, a service edge **112** and any other intervening elements. As illustrated in FIG. **3**, the carrier-tagged flow payload **422** may be kept substantially intact, and a carrier tag **420** may be prepended to the carrier-tagged flow payload **422** to prepare it for transmission through the access network. Depending on implementation, the carrier tag **420** may include, for example, a pseudowire label, a VLAN identifier, a tunnel label or the like. Multiple carrier tags may be stacked within a message or frame to provide for a hierarchical aggregation and routing mechanism to be implemented in the access network.

In accordance with one exemplary embodiment, the building aggregator system **114**, such as via interface modules **310-316**, may perform the aforementioned carrier tagging of traffic received from CPE **116** and bound for the service edge **112**. Likewise, the building aggregator system may receive

from the service edge messages having such carrier tags and will remove the carrier tags and distribute the messages to the appropriate CPE-side link.

It is particularly noteworthy in FIG. **3** that, regardless of message type, all of the carrier tags **420** may be of uniform format. (In the case of tunnel labels, for example, messages of different types may even have the same tag value if they happen to be commonly routed.) The use of a uniform carrier tag format for all message types makes it possible for simple, generic handling of all traffic types through the access network using a uniform set of network elements that process traffic based on carrier tags. The switching elements within the access network may simply inspect the carrier tag(s) **420** of messages to determine how the message should be switched or routed without regard to message type or contents. In this manner, the access network becomes "service agnostic" and does not have to be concerned with the specifics of the protocols or addressing imbedded in the customer traffic. The generic nature of the carrier tag also allows for readily supporting any other message types not shown in FIG. **3**, with little or no changes being required in the design and operation of the layer-2 switches **118** or other elements.

In some implementations, it may be desirable to prepend one or more tunnel labels (not shown) to the messages **410-418**. This may be in addition to any pseudowire labels or other carrier tags already applied. A tunnel label allows a tunnel to be established throughout the access network, such as between a building aggregator and a service edge, improving scalability in the network. This mechanism may be particularly useful when many flows are to be routed to the same destination or service edge. By assigning the flows to a common tunnel, network elements, such as the layer-2 switch **118**, may collectively route the service emulation instances within the tunnel by evaluating the tunnel label. In an exemplary embodiment, the tunnel label may be an LSP label prepended to the messages **410-418**. In accordance with exemplary embodiments, tunnel labels may also be stacked to any degree needed to support a tunneling hierarchy, which may further facilitate efficient and scalable management of large numbers of flows.

Although the carrier-tagged flow payload **422** is shown and described as being kept essentially intact, it may be desirable in some situations to modify this original message. For example, the original message portion **422** of the Ethernet frame message **410** and the frame relay frame **414** may include a frame check sequence (FCS). In many networks, the FCS is not used and may be removed. In other cases, the Ethernet frame check sequence (FCS) as received in the Ethernet frame may optionally be included, as is, rather than being deleted or recalculated by the building aggregation system **114**. This can be advantageous for detecting errors or corruption of the data that might occur as the customer payload traverses the network.

FIG. **4** illustrates some of the functions that may be performed by the building aggregation system **114** in accordance with one exemplary embodiment. The functions described herein may be performed by software, hardware, or some combination thereof. Furthermore, the functions may be performed by one of the elements or a combination of the elements of the building aggregation system discussed above with reference to FIG. **4**.

For traffic originating at a CPE **116** and bound for a service edge **112**, a building aggregation system **114** may receive messages from multiple customer CPEs, where the messages may be received along different physical or logical ports associated with each customer flow. The building aggregation system may append one or more carrier tags to the incoming messages and otherwise encapsulate or process the messages to prepare them for transmission through the access network **100**. The building aggregation system may send the carrier-tagged messages into the access network, which may be in the form of an aggregated flow over a wideband or broadband communications link. Likewise, the building aggregation system may operate in a complementary manner for traffic coming from a service edge that is bound for a CPE. The building aggregation system may receive carrier-tagged messages that have traversed the access network, interpret the carrier tags to determine where to distribute the message, 'decapsulate' the messages to remove the carrier tags, and distribute the messages to the appropriate physical and logical ports leading to the CPE.

To properly route traffic in either direction between each of one or more CPE flows and a corresponding access tunnel, the service emulation instance or logical networking tagged path, the building aggregation system **114** may maintain an association between a customer interface port and a particular carrier tag value or set of values. Where service emulation instances are used in the access network **100**, the building aggregation system may maintain an association between customer-facing ports and a service emulation instance mapping identifier. Where pseudowires are used in the access network, the building aggregation system may maintain an association between customer-facing ports and a pseudowire label value. One or more tunnel labels or LSP labels may also be associated with customer ports, either directly or indirectly by association with pseudowires or service emulation instances. Any of these associations may be maintained, for example, in a look-up table which may be referred to as a management information base (MIB) or a forwarding information base (FIB). These associations may be established by negotiation or coordination with other access networks elements or by control applied to the building aggregation system by a provisioning system. The building aggregation system may also perform functions related to controlling traffic volumes and related to remotely managing demarcation devices.

As depicted in FIG. **4**, the functions of a building aggregation system **114** may include a policing function **510**, a packing/unpacking function **512**, a maximum transmission unit (MTU) control **514**, a demarcation device manager **516**, a provisioning and operations function **518**, and a control and maintenance function **520**. Exemplary embodiments are more particularly described below with respect to the policing function. For more information as to the other functions, reference is made to U.S. patent application Ser. No. 10/858,503, entitled: Method and Apparatus for Processing Labeled Flows in a Communications Access Network, the content of which is hereby incorporated by reference in its entirety.

The policing function **510** may be responsible for monitoring a customer's traffic and controlling the customer's bandwidth, or rate at which customer traffic enters and passes through the access network **100**. The policing function may involve marking or dropping of customer packets or frames that exceed configured threshold values. Dropping or discarding of packets may occur when traffic from the customer exceeds a maximum bandwidth. Packets that represent an excessive bandwidth may simply not be forwarded to the access network. When this happens, the policing function may notify the respective CPE **116** of the data loss, whereupon the CPE may throttle back or take action to recover from the lost data. The building aggregation system may also engage in customary flow control coordination with CPE to regulate incoming bandwidths of flows, for example, using

"pause" or "ready to receive" control messages of some nature that are well known among those of ordinary skill in the art.

In contrast to dropping packets, marking relates to designating which packets, if any, may be dropped later during transmission if the network or a network element becomes too busy to handle all of the customer's traffic. Generally, a customer may contract with a service provider for a guaranteed minimum bandwidth, as well as an absolute maximum or "burstable" bandwidth that may be supported as long as the network is not too congested to handle it. At times, the customer may attempt to send traffic at a bandwidth in excess of the guaranteed bandwidth. While the network may generally accommodate the higher bandwidth in these situations, it is desirable to set a discard eligibility bit in the message header. This allows downstream network elements to selectively discard some of the customer's traffic if they are overburdened with other network traffic. In accordance with one exemplary embodiment, the least significant bit (LSB) within an Experimental field of an MPLS label maybe used to indicate discard eligibility. In effect, when a customer has exceeded a guaranteed bandwidth and begins to experience data loss due to network congestion, the customer may reduce bandwidth or take other measures until the network congestion subsides. In some cases, transport protocols operating within the CPE **116** may sense the data loss and may automatically perform flow control to effectively reduce the bandwidth. A marking process that may be performed as part of policing function **510** is discussed in greater detail below with reference to FIG. **5**. Policing function **510** may also encompass metering or rate limiting to adjust the timing of forwarding of packets for better control of variable rate traffic.

In accordance with one exemplary embodiment, policing of customer traffic may extend across multiple service types and multiple logical or physical ports. For example, a customer may contract for two individual flows having certain committed minimum and maximum burstable bandwidths. The QoS measures implemented in the building aggregation system **114** may also control the sum of the two flows to remain below a certain limit, even if the flows are of different types or enter the building aggregation system along different ports. Minimum and maximum bandwidths may even be applied in a hierarchical nature, for example, by company, department, employee, device, service type or specific flow.

FIG. **5** is a flowchart including various steps in a method of marking or dropping of customer packets that may be performed by the policing function **510** (see, e.g., FIG. **4**) of the building aggregation system **114** in accordance with an exemplary embodiment. The process begins in step **610**, wherein a unit of data traffic, such as a data frame, is received from CPE **116**. As discussed above, the data may be received from a variety of CPE, such as an Ethernet client **116***a*, a private line customer **116***b*, a frame relay **116***c* an ATM **116***d* or the like.

In step **612**, a service profile for the customer is retrieved from memory or other data storage devices. The service profile describes attributes of services to which the customer is subscribed. Each service may have its own requirements, such as a guaranteed or "committed" minimum data rate that the service provider promises to make available at all times. The guaranteed amount of bandwidth as specified in the service profile is compared to the bandwidth currently being consumed by the customer in step **614**.

If a determination is made in step **614** that the data is in excess of the guaranteed data rate as specified in the service profile, then execution proceeds to step **616**, wherein a discard eligibility bit is set to indicate to downstream network elements that the data frame may be regarded as discardable if the network is busy or heavily congested. In accordance with one exemplary embodiment, a discard eligibility indication may even be applied to traffic types which lack an equivalent attribute in the native protocol. For example, the least significant bit (LSB) within the Experimental field of an MPLS label maybe used to indicate discard eligibility. Thus, consistent QoS measures may be extended to flows of all types.

Those of ordinary skill in the art will recognize how steps **614** and **616** may actually be implemented using a "token bucket" approach or other techniques to discern data frames that represent excess beyond a certain volume per unit time. Policing or rate limiting of customer flows may also take into account a maximum burst duration that may be configured for a flow according to a service profile that describes the flow attributes.

Whether or not step **616** is performed responsive to the determination in step **614**, processing proceeds to step **618**, wherein the data is prepared and transmitted to the service edge **112** via the associated carrier-tagged flow. Preparation of the data for transmission may include, for example, formatting the data for transport via a TDM communications link or a packet-based communications link, identifying the correct carrier-tagged flow that is to carry the data associated with the data, appending a carrier tag to the data, and other processing as described herein.

In accordance with exemplary embodiments, an apparatus, method and computer-readable storage medium may be provided whereby a customer, service provider, or other entity in communication with the customer or service provider may calculate an estimated throughput requirement of the customer, such as to facilitate selecting the minimum bandwidth, maximum bandwidth or some other appropriate bandwidth. In accordance with exemplary embodiments, the estimated throughput requirement of the customer may be calculated based on an estimated bandwidth requirement of the customer. In addition to the bandwidth requirement, the estimated throughput requirement may be calculated further based on the size of the customer's packets and overhead in sending the customer's packets across one or more core networks, and the line rate for sending the customer's packets across the core network(s).

As shown and described herein, the throughput requirement may be a maximum throughput requirement calculated based on a maximum bandwidth. It should be understood, however, that exemplary embodiments may be equally applicable to a minimum throughput requirement calculated based on a minimum bandwidth. Further, the estimated throughput requirement may be calculated with respect to one or more customer flows including tunnel labels (e.g., LSP labels) and service emulation identifiers (e.g., pseudowire labels); and transmitted utilizing Ethernet as the layer-2 protocol, and thus encapsulated in an Ethernet frame. Further, in various instances of this exemplary embodiment, a logical networking tagged flow may be used such that one or more customer flows may further include an outer VLAN tag. Again, however, it should be understood that exemplary embodiments may be equally applicable to customer flows formatted in a number of other manners.

Referring now to FIG. **6**, a method of calculating a throughput requirement of a network, or more particularly a customer accessing one or more network services, in accordance with one exemplary embodiment begins in step **710**, wherein an estimated bandwidth requirement is received, such as from the customer. One exemplary technique for calculating or otherwise selecting an estimated bandwidth requirement of a customer is disclosed in U.S. patent application Ser. No.

11/873,827 entitled: System, Method and Computer-Readable Storage Medium for Calculating Addressing and Bandwidth Requirements of a Network, filed concurrently herewith, now U.S. Pat. No. 8,014,316, the content of which is hereby incorporated by reference in its entirety. The estimated bandwidth requirement may be received in any of a number of different formats, but in one exemplary embodiment are received in megabits per second (Mbps) and bytes, respectively.

In step **712**, a layer-2 frame rate requirement may be calculated based on the estimated bandwidth requirement and a frame size. In this regard, the frame size may, for example, be received in a manner similar to that of the estimated bandwidth requirement. Alternatively, for example, the frame size may be predetermined or selected from a predetermined list of different frame sizes. Similar to the bandwidth requirement, the frame rate requirement may be calculated in any of a number of different formats, but in one exemplary embodiment, may be calculated in frames per second (fps). The fps requirement may be calculated based on the estimated bandwidth requirement accounting for the frame size, and further accounting for the gap between the transmission of frames of data (inter-packet gap) and overhead in transmitting a frame (e.g., preamble, start-of-frame delimiter). The core network(s) may have a maximum, or line, rate for the transmission of data packets, which may take into account overhead required for the tunnel control plane. And as such, the fps requirement may be limited by that line rate, or rather throughput at a bandwidth of that line rate (accounting for the frame size, inter-packet gap and overhead, and further accounting for the tunnel label, the service emulation identifier, and the Ethernet frame within which the customer's frame is encapsulated for transmission across the core network(s)).

Written notationally, the frame rate requirement (L2_fs_req) may be calculated as follows:

$$\text{L2_fps_req} = \text{MIN}\left[\left(\frac{\frac{BW \times 1000000}{8}}{G + OH + LN + P}\right), \frac{LR}{G + OH + LN + TL + SEI + EF + P}\right] \quad (1)$$

In the preceding equation (1), BW represents the estimated bandwidth requirement in Mbps (multiplication by 1000000 and division by 8 thereby converting the bandwidth from Mbps to bytes). In this regard, the estimated bandwidth requirement may represent a maximum bandwidth requirement, a minimum bandwidth requirement or another appropriate bandwidth requirement. Also in the preceding, G, OH, LN and P represent the inter-packet gap, overhead size (e.g., preamble, start-of-frame delimiter, etc.), any applicable outer logical networking (e.g., VLAN) tag size, and packet size, respectively, in bytes. Further, LR, TL, SEI and EF represent the line rate, tunnel label size, service emulation identifier size and Ethernet frame size, respectively, in bytes.

The inter-packet gap, overhead size, packet size, line rate, tunnel label size, service emulation identifier size and Ethernet frame size may be predefined or otherwise known. Thus, in a more particular example, the fps requirement may be calculated as follows:

$$\text{L2_fps_req} = \text{MIN}\left[\left(\frac{\frac{BW \times 1000000}{8}}{20 + 4 + P}\right), \frac{LR}{20 + 4 + 12 + 22 + P}\right] \quad (2)$$

where if BW>1000 Mbps, LR=1249375000 bytes; else, LR=124937500 bytes. Also in equation (2), the value 20 represents the inter-packet gap (G) and overhead size (OH) in bytes, and 4 represents the size of an outer VLAN tag (although this may not be included in instances where a logical networking tagged flow is not used). Further, the value 12 represents the tunnel label (TL) and service emulation identifier (SEI) sizes in bytes, and 22 represents the Ethernet frame size (EF) in bytes.

In step **714**, a layer-2 throughput requirement may be calculated based on the calculated layer-2 frame rate (e.g., fps) requirement. Similar to the other values, the throughput requirement may be calculated in any of a number of different formats, but in one exemplary embodiment, may be calculated in bits per second (bps) or Mbps. The throughput requirement may be calculated based on the calculated layer-two frames per timer period requirement, and further based on the frame size. Written notationally, the throughput requirement (L2_bps_req) may be calculated in bits as follows:

$$L2_bps_req = 8 \times L2_fps_req \times P \quad (3)$$

where the value 8 converts the frame size from bytes to bits in instances in which the throughput requirement is represented in bits. Thus, given the example values in equation (2), the frame rate and throughput requirements for an estimated bandwidth requirement of 100 Mbps (i.e., BW=100 Mbps) may be calculated as shown in Table 1 below (for the case using a logical networking tagged flow).

TABLE 1

Customer L2 Packet Size (P)	L2_fps_req	L2_bps_req
64	142,045	72,727,273
78	122,549	76,470,588
128	82,237	84,210,526
256	44,643	91,428,571
404 (Internet Mix)	29,206	94,392,523
512	23,321	95,522,388
1024	11,927	97,709,924
1522	8,085	98,447,607

The access network **100** of exemplary embodiments may be utilized in any of a number of different configurations to provide any of a number of different types of services to one or more customers. Examples of such configurations and services include, for example, metropolitan, national or international Ethernet virtual private line (EVPL), Ethernet access (EA) to Layer-3 Virtual Private Network (VPN), Layer-2 EA to Layer-2 VPN, and virtual private LAN service (VPLS). Also, in various instances, a customer may have one or more sites **110**, each of which may use one or more services having one or more throughput requirements. And in such instances, a service provider may offer a converged access for one or more services. Thus, in accordance with exemplary embodiments, the customer, service provider, or other entity in communication with the customer or service provider may further calculate an estimated converged throughput requirement for a customer's site including one or more traffic flows of one or more services having one or more throughput requirements.

Written notationally, an estimated converged throughput requirement (Conv_L2_bps_req) for a customer's site **110** may be calculated as follows:

$$\text{Conv_L2_bps_req} = \sum_{i=1}^{I} \text{L2_bps_req}_i \times n_{\text{L2_bps_req}_i} \quad (4)$$

In the preceding, i=1, . . . I represents unique service/throughput requirements for the customer's site (separately accounting for the same service with different throughput requirements, as well as different services with the same or different throughput requirements), and $n_{L2_bps_req}$, represents the number of flows with the respective service/throughput requirement.

Further, in various instances the service provider may offer different types of converged service, where the types of converged service may differ from another in a number of different manners such as whether the services are offered in an on-network environment or an off-network environment. Further, the service provider may offer one or more variations for each type of converged service. More particularly, for example, the service provider may offer various billing platform options for each type of converged service. And within a type and billing platform option, the service provider may make available incremental bandwidths to customers, and may even offer various increments with some types/options and not others. In these instances, calculation of the estimated converged throughput requirement for a customer's site may further account for the type of converged service, and billing platform option with which that type of service is being (or will be) provided to the respective customer.

For purposes of illustration, consider an example where on-network converged access and off-network converged access may be offered with different billing platform options at the following bandwidths in Table 2:

TABLE 2

	On-Network		Off-Network	
BW (Mbps)	Option One	Option Two	Option One	Option Two
1	X	X	X	X
2	X	X	X	X
3	X	X	X	X
4	X	X	X	X
5	X	X	X	X
6	X	X	X	X
7	X	X	X	X
8	X	X	X	
9	X	X	X	
10	X	X	X	
15	X	X		
20	X	X	X	
25	X	X		
30	X	X	X	
35	X	X		
40	X	X	X	X
45	X	X		
50	X	X	X	
60	X	X	X	
70	X	X	X	
80	X	X	X	
90	X	X	X	
100	X	X	X	
150	X	X	X	X
200	X	X	X	
250	X	X		
300	X	X	X	
350	X	X		
400	X	X	X	
450	X	X		
500	X	X	X	
600	X	X	X	X
700	X	X	X	
800	X	X	X	
900	X	X	X	
1,000	X	X	X	X

In instances such as that of Table 2, the service provider's offered converged throughput ($\text{Conv_L2_bps_on}_{1,2}$) for selection by a customer in an on-network environment may be calculated for both billing platform options one and two, and based on the customer's converged throughput requirement (Conv_L2_bps_req) as follows:

If Conv_L2_bps_req<11 Mbps,
$\text{Conv_L2_bps_on}_{1,2}$=CEILING (Conv_L2_bps_req, 1);
else,
If Conv_L2_bps_req<51 Mbps,
$\text{Conv_L2_bps_on}_{1,2}$=CEILING (Conv_L2_bps_req, 5);
else,
If Conv_L2_bps_req<101 Mbps,
$\text{Conv_L2_bps_on}_{1,2}$=CEILING (Conv_L2_bps_req, 10);
else,
If Conv_L2_bps_req<501 Mbps,
$\text{Conv_L2_bps_on}_{1,2}$=CEILING (Conv_L2_bps_req, 50);
else,
$\text{Conv_L2_bps_on}_{1,2}$=CEILING (Conv_L2_bps_req, 100).

In the preceding, CEILING (x, y) is a function for rounding a number x to the nearest multiple of significance y.

Continuing the example above, the service provider's offered converged throughput (Conv_L2_bps_off_1) for selection by a customer in an off-network environment may be calculated for billing platform option one as follows:

If Conv_L2_bps_req<11,
Conv_L2_bps_off_1=CEILING (Conv_L2_bps_req, 1);
else,
If Conv_L2_bps_req<101,
Conv_L2_bps_off_1=CEILING (Conv_L2_bps_req, 10);
else,
If Conv_L2_bps_req<151,
Conv_L2_bps_off_1=CEILING (Conv_L2_bps_req, 50);
else,
Conv_L2_bps_off_1=CEILING (Conv_L2_bps_req, 100).

And the service provider's offered converged throughput (Conv_L2_bps_off_2) for selection by a customer in an off-network environment may be calculated for billing platform option two as follows:

If Conv_L2_bps_req<8,
Conv_L2_bps_off_2 CEILING (Conv_L2_bps_req, 1);
else,
If Conv_L2_bps_req<41,
Conv_L2_bps_off_2=CEILING (Conv_L2_bps_req, 40);
else,
If Conv_L2_bps_req<151,
Conv_L2_bps_off_2=CEILING (Conv_L2_bps_req, 150);
else,
IF Conv_L2_bps_req<601,
Conv_L2_bps_off_2=CEILING (Conv_L2_bps_req, 600);
else,
Conv_L2_bps_off_2=CEILING (Conv_L2_bps_req, 1000).

One or more embodiments may be implemented as a method, a device, or a computer program product. Accordingly, an embodiment may take the form of an entirely hardware embodiment, an entirely software embodiment, or an embodiment combining software and hardware aspects. Furthermore, implementations of an embodiment may take the form of a computer program product including a computer-readable storage medium having computer-readable program instructions (e.g., computer software) embodied in the storage medium. More particularly, implementations of certain embodiments may take the form of one or more computer software macros operable with spreadsheet computer software such as Excel, distributed by the Microsoft Corporation of Redmond, Wash. Any suitable computer-readable storage medium may be utilized including hard disks, CD-ROMs, optical storage devices, or magnetic storage devices.

In the preceding specification, various embodiments have been described. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims that follow. The specification and drawings are accordingly to be regarded in an illustrative rather than restrictive sense.

What is claimed is:

1. An apparatus for establishing communication according to a throughput requirement of a network, the network comprising a building aggregation system, a service edge, and a layer 2 switch between the building aggregation system and the service edge, the apparatus comprising:

a processor configured to

receive a bandwidth requirement from Customer Premise Equipment (CPE) coupled to the building aggregation system,

calculate a data link-layer frame rate requirement for the layer 2 switch based on the bandwidth requirement and an Ethernet frame size, the Ethernet frame size selected from a predetermined list of different frame sizes, and calculating the data link-layer frame rate requirement including limiting the calculated data link-layer frame rate requirement by a line rate for transmission of data packets across a core network,

calculate a data link-layer throughput requirement based on the calculated data link-layer frame rate requirement and the Ethernet frame size, and

establish communication, over the network, between the service edge and the CPE according to the data link-layer throughput requirement;

wherein the data link-layer frame rate requirement, L2_fps_req is calculated as follows:

$$\text{L2_fps_req} = \text{MIN}\left[\left(\frac{\frac{BW\times 1000000}{8}}{G+OH+LN+P}\right), \frac{LR}{\frac{G+OH+LN+TL+}{SEI+EF+P}}\right]$$

wherein BW represents the bandwidth requirement in Mbps, G, OH, LN, and P represent an inter-packet gap, overhead size, any applicable outer logical networking tag size, and packet size, respectively, in bytes; and

wherein LR, TL, SEI, and EF represent line rate, tunnel label size, service emulation identifier size, and the Ethernet frame size, respectively, in bytes.

2. The apparatus according to claim **1**, wherein the processor is configured to calculate one or more data link-layer throughput requirements for one or more traffic flows of one or more network services having one or more throughput requirements, and

wherein the processor is further configured to calculate a converged throughput requirement based on the one or more data link-layer throughput requirements.

3. The apparatus according to claim **2**, wherein the processor is configured to calculate the converged throughput requirement accounting for a type of converged access, different types of converged access having different available bandwidths.

4. The apparatus according to claim **3**, wherein the processor is configured to calculate the converged throughput requirement accounting for a billing platform for a type of converged access, different billing platforms for a type of converged access having different available bandwidths.

5. A method of establishing communication according to a throughput requirement of a network, the network comprising a building aggregation system, a service edge, and a layer 2 switch between the building aggregation system and the service edge, the method comprising:

receiving, by a processor, a bandwidth requirement from Customer Premise Equipment (CPE) coupled to the building aggregation system;

calculating, by the processor, a data link-layer frame rate requirement for the layer 2 switch based on the bandwidth requirement and an Ethernet frame size,

calculating, by the processor, the data link-layer frame rate requirement including limiting the calculated data link-layer frame rate requirement by a line rate for transmission of data packets across a core network; and

calculating, by the processor, a data link-layer throughput requirement for the network based on the calculated data link-laver frame rate requirement and the frame size, and

establishing, by the processor, communication between the service edge and the CPE according to the data link-layer throughput requirement;

wherein the data link-layer frame rate requirement, L2_fps_req is calculated as follows:

$$\text{L2_fps_req} = \text{MIN}\left[\left(\frac{\frac{BW\times 1000000}{8}}{G+OH+LN+P}\right), \frac{LR}{\frac{G+OH+LN+TL+}{SEI+EF+P}}\right]$$

wherein BW represents the bandwidth requirement in Mbps, G, OH, LN, and P represent an inter-packet gap, overhead size, any applicable outer logical networking tag size, and packet size, respectively, in bytes; and

wherein LR, TL, SEI, and EF represent line rate, tunnel label size, service emulation identifier size, and Ethernet frame size, respectively, in bytes.

6. The method according to claim **5**, wherein calculating a data link-layer throughput requirement comprises calculating one or more data link-layer throughput requirements for one or more traffic flows of one or more network services having one or more throughput requirements, and

wherein the method further comprises calculating a converged throughput requirement based on the one or more data link-layer throughput requirements.

7. The method according to claim 6, wherein calculating a converged throughput requirement comprises calculating a converged throughput requirement accounting for a type of converged access, different types of converged access having different available bandwidths.

8. The method according to claim 7, wherein calculating a converged throughput requirement comprises calculating a converged throughput requirement accounting for a billing platform for a type of converged access, different billing platforms for a type of converged access having different available bandwidths.

9. A non-transitory computer-readable storage medium having computer-readable program code portions stored therein for establishing communication according to a throughput requirement of a network calculated via a processor configured to read the program code portions, the network comprising a building aggregation system, a service edge, and a layer 2 switch between the building aggregation system and the service edge, the computer-readable program code portions comprising:

a first executable portion configured to cause the processor to receive a bandwidth requirement from Customer Premise Equipment (CPE) coupled to the building aggregation system;

a second executable portion configured to cause the processor to calculate a data link-layer frame rate requirement for the layer 2 switch based on the bandwidth requirement and an Ethernet frame size, the Ethernet frame size selected from a predetermined list of different frame sizes, and to calculate the data link-layer frame rate requirement including limiting the calculated data link-layer frame rate requirement by a line rate for transmission of data packets across a core network;

a third executable portion configured to cause the processor to calculate a data link-layer throughput requirement based on the calculated data link-layer frame rate requirement and the Ethernet frame size; and

a fourth executable portion configured to cause the processor to establish communication, over the network, between the service edge and the CPR according to the data link-layer throughput requirement;

wherein the data link-layer frame rate requirement, L2_fps_req is calculated as follows:

$$\text{L2_fps_req} = \text{MIN}\left[\left(\frac{\frac{BW \times 1000000}{8}}{G + OH + LN + P}\right), \frac{LR}{G + OH + LN + TL + SEI + EF + P}\right]$$

wherein BW represents the bandwidth requirement in Mbps, G, OH, LN, and P represent an inter-packet gap, overhead size, any applicable outer logical networking tag size, and packet size, respectively, in bytes; and

wherein LR, TL, SEI, and EF represent line rate, tunnel label size, service emulation identifier size, and Ethernet frame size, respectively, in bytes.

10. The non-transitory computer-readable storage medium according to claim 9, wherein the third executable portion is configured to calculate one or more data link-layer throughput requirements for one or more traffic flows of one or more network services having one or more throughput requirements, and wherein the computer-readable program code portions further comprise a fourth executable portion configured to calculate a converged throughput requirement based on the one or more data link-layer throughput requirements.

11. The non-transitory computer-readable storage medium according to claim 10, wherein the fourth executable portion is configured to calculate the converged throughput requirement accounting for a type of converged access, different types of converged access having different available bandwidths.

12. The non-transitory computer-readable storage medium according to claim 11, wherein the fourth executable portion is configured to calculate the converged throughput requirement accounting for a billing platform for a type of converged access, different billing platforms for a type of converged access having different available bandwidths.

* * * * *